Leveraging Installed Copper to Reach Underserved and Unserved Community Anchor Institutions

Hatteras Networks

June 1, 2009

About Hatteras Networks

Hatteras Networks has redefined Ethernet services, unleashing them from fiber and opening them up to the majority of business locations. Hatteras Networks' EtherFLEX service delivery solutions make Ethernet services available over copper, allowing Service Providers worldwide to expand their Metro Ethernet market -- previously been limited to the fiber footprint. Historically, T1s and E1s have been the fundamental building blocks for voice and data business services. Hatteras' solution enables Service Providers to cost- effectively offer services at rates over 50 Mbps using existing copper facilities. This innovation allows businesses to migrate from legacy Frame Relay, ATM and T1/E1 connections to transparent Ethernet services for voice and data business connectivity, infrastructure backhaul and mobile wireless backhaul solutions. Already deployed on five continents, Hatteras Networks solutions are enabling a market expansion for carriers that are as significant for business services as "triple play" is for residential.

Contents

Executive Summary

Introduction: The broadband funding provisions (Broadband Stimulus) of the American Reinvestment and Recovery Act (ARRA) create the opportunity to provide broadband access to unserved and underserved communities. While the Broadband Stimulus effort is designed to address these challenges and enable broadband to anchor tenants in these environments, solutions for provisioning the last mile to accommodate the explosive demand for broadband business access are quite limited.

Ethernet is the obvious choice to meet the needs of businesses and governmental organizations. Ethernet is dominant in enterprise networks; it has a very lost "cost per bit." Ethernet is highly flexible. It can be delivered on nearly all copper, optical fiber and wireless physical media, and its bandwidth has continued to increase. Ethernet's dominance is also the result of its ease of deployment and ease of use. Given its ubiquity, end users are increasingly demanding wide-area connectivity over Ethernet interfaces.

There are three primary media – fiber, wireless, and copper - that can be leveraged to deliver Ethernet-based broadband access. Each of these media has its advantages and disadvantages. Despite the hype around fiber deployments, and despite its appeal as part of the Broadband Stimulus that it is "shovel-ready", fiber deployments are not economically viable for the vast majority of un/underserved anchor tenants, leaving many entities out of reach. The best solution will take advantage of all three media, but copper-based Ethernet access (Ethernet-over-Copper) must be included in order to cost effectively reach all critical organizations.

Fiber deployments have received a lot of attention over the past few years, and, as a result, many of us have been under the impression that the Fiber access market is larger than the copper access market. This misconception is understandable, but it is a misconception. Copper reaches over 5 times the number of commercial buildings and mobile towers than fiber reaches, and fiber deployment is growing at a snail's pace – increasing at less than 2 percentage basis point per year on average for the past 5 years.

The expansive wealth of copper in the ground that has been deployed by large and small telephone providers alike over the last 50 years and paid for by the ratepayers is a key nationwide and ubiquitous asset. Ethernet-over-Copper takes advantage of this asset, and in today's financial market environment it is essential that solutions implemented as part of the Broadband Stimulus take advantage of all the assets that are readily and economically available. The advantages to the Broadband Stimulus efforts in taking advantage of this installed copper include the ability to provision:

- Greater than 10 anchor tenants with copper for the price of provisioning one anchor tenant with fiber,
- Provisioning anchor tenants in intervals at least 10 times faster with copper than with fiber, and
- Providing virtually all anchor tenants with 10Mbps, 20Mbps or even 50Mbps Ethernet, regardless of whether fiber can be pulled to those locations. Today optical networks only reach about 20% of the buildings in the US[1], and most of those are in urban areas.

[1] According to Vertical Systems Group (www.verticalsystems.com).

The Services: Ethernet-over-Copper (EOC) cost effectively fills the large gap in the service reach for business-grade solutions, which are critical to rural area economic development. It provides a seamless provising of metro Ethernet to the more than 80% of buildings that are not reachable by fiber[2], including SMB and enterprises, schools and other educational institutions, health care facilities, and other government locations. Ethernet-over-Copper offers these anchor institutions a familiar service demarcation point, delivers symmetrical data rates typically targeted in the 2-50 Mbps range, and provide support for the full range of carrier-class Ethernet service features such as multiple Ethernet Virtual Circuits (VLANs) per customer, bandwidth profiles, VLAN switching, and Quality of Service (QOS) performance profiles to guarantee service quality.

The Applications: Ethernet-over-Copper supports in a variety of applications, including business access, cellular backhaul, WiFi backhaul, DSLAM backhaul, and in-building access.

- <u>Regional economic development</u>: In addition to supporting existing anchor institutions, EoC provides the high bandwidth, symmetric, and cost effective last mile Ethernet access required to attract new small and medium businesses and large enterprise businesses to the target un/underserved areas. EoC provides a wealth of Ethernet service types required to support these organizations including point-to-point E-Line service [a service connecting two customer Ethernet ports over a WAN.], multipoint-to-multipoint transparent E-LAN service [a multipoint service connecting a set of customer endpoints, giving the appearance to the customer of a bridged Ethernet network connecting the sites], and high-speed commercial Internet access. These services are required to support the demands of businesses, local health care providers, first responders, government, and educational entities, which require guaranteed service (strict service level agreements). Furthermore, because the services provide full VLAN support, multiple end-to-end connections may be delivered over the same Ethernet-over-Copper service further increasing the cost effectiveness of the solution in un/underserved areas.

- <u>Wireless Enablement</u>: With the growth of mobile voice and data services, fixed wireless services based on WiFi and WiMax, and DSLAM based services, capacity requirements from the base station or access point to the nearest switching POP are expanding. EoC directly enables not only the delivery of wired services, but also enables next generation mobile and wireless services to targeted areas by providing a cost effective high speed backhaul from rural mobile tower locations. Delivering the increased bandwidth needed to enable today's and tomorrows advanced and evolving mobile / wireless applications to rural subscribers is a critical part of addressing the needs of the underserved areas. EoC provides a cost effective, resilient, rapidly deployable alternative to the cost and complexity of laying fiber to extend high bandwidth services to these rural tower locations.

- <u>DSL Expansion</u>: Today rural DSL is limited in its reach and service scope due to the rural providers' ability to provide sufficient bandwidth out to the remote cabinet or DSLAM device. As discussed later, EoC provides a cost effective, resilient, and

[2] According to Vertical Systems Group (www.verticalsystems.com).

rapidly deployable alternative to the cost and complexity of laying fiber to these remote serving units, broadening the availability of high bandwidth services and increased service quality in rural target markets.

The Technology: Ethernet-over-Copper employs Metro Ethernet Service Edge technologies based on two technologies:

- Mid-Band Ethernet - the IEEE 802.3ah Ethernet in the First Mile standard ratified in 2004. The long reach option of 802.3ah used for Ethernet-over-Copper, known as 2BASE-TL, defines a multi-pair bonding scheme, which is transparent to the Ethernet layer. Using this bonding technique across up to 8 pairs, Ethernet services can be offered with data rates of:
    - 10 Mbps up to about 15,000 feet. Extendable to 120,000 ft with repeaters.
    - 20 Mbps up to about 10,000 feet. Extendable to 80,000 ft with repeaters.
    - 45 Mbps up to about 5,000 feet. Extendable to 40,000 ft with repeaters.

  The key benefit from using Mid-Band Ethernet is its high cost effectiveness and bandwidth capabilities.

- Ethernet-over-TDM (EoTDM) - T1 at 1.544 Mbps has been a key access technology for business voice and data services for decades. From its beginnings as a voice trunk line technology to Internet access for small and medium sized businesses, T1s have proven to be a well-understood and versatile last mile technology. These lines reach nearly every business and Ethernet can be transported over T1 as a single link or bonded group of links allowing service providers to deliver Ethernet at rates from 1 Mbps up to 16 Mbps. There are three standards methods for delivering Ethernet over T1 lines. These are multilink point-to-point protocol (MLPPP, specified in IETF RFC 1990 and RFC 3518), GFP/VCAT (specified in ITU-T G.7041 and G.7043) and G.bond or EFM (specified by ITU-T G.992.2). While each technology has its strengths, they all deliver comparable performance and are available from multiple equipment vendors. The key benefit from using Ethernet-over-NxT1 for delivering Ethernet services is that the service provider is able to reach all of their customer locations, regardless of geography and proximity to their facilities. T1 facilities are not distance impacted because the service provider network uses repeaters that regenerate the transmission and extend the network.

The bonding used both for Mid-Band Ethernet and Ethernet-over-NxT1 brings the additional benefit of resiliency – a feature demanded by many enterprise and government end users. Because there are multiple links involved in the access method, it is inherently protected against one or more of those links' being interrupted – for example by a backhoe or an excavator.

Background and Introduction

Pulled by enterprise demand for higher capacity packet data services and pushed by its ubiquity, flexibility and low cost structure, Ethernet has become the accepted protocol for packet transmission for broadband delivery for building and campus local area networks. However, Service Providers have had no viable options for delivering broadband Ethernet services to non fiber-fed business sites.

Ethernet-over-Copper cost effectively fills the gap in the service portfolio and enables all of the benefits of fiber-based Ethernet services to be offered to the vast majority of customer locations that do not have access to fiber.

What is Ethernet-over-Copper?
Ethernet-over-Copper (EoC) provides a seamless provisioning of metro Ethernet to smaller and rural customer locations over existing copper facilities. EoC offers: a familiar 10/100Base-T RJ-45 jack as a service demarcation point; symmetrical services typically targeted in the 2-20 Mbps range (although service capacity may be provisioned over 45Mbps); and can offer support for the full range of carrier-class features such as multiple Ethernet Virtual Circuits (i.e. VLANs), VLAN switching, bandwidth profiles, and Quality of Service (QOS) performance profiles as defined by the Metro Ethernet Forum (MEF).

With Ethernet-over-Copper, end-to-end Ethernet services become universal – customers can now be reached with optical or copper access.

How Ethernet-over-Copper is Delivered
Ethernet-over-Copper employs Metro Ethernet Service Edge technologies based on the IEEE 802.3ah Ethernet in the First Mile standard ratified in 2004. The long reach option of IEEE 802.3ah used for Ethernet-over-Copper, known as 2BASE-TL, was designed as native Ethernet technology, with implementations providing nominal symmetric data rates of 2.3 Mbps over a single pair of voice-grade copper over CSA (carrier serving area) distances of 2700-3600m (9-12Kft) to a maximum data rate of 15.4 Mbps over shorter pairs (up to 1300ft).

For a more in-depth explanation of the operation of the underlying Ethernet-over-Copper technology, please see the Ethernet-over-Copper Technology Handbook: IEEE 802.3ah 2BASE-TL for Executives available from Hatteras Networks found in Appendix A.

Availability, Simplicity and Consistency
The most significant benefit of Ethernet-over-Copper is the ability to readily serve the over 80% of business locations not currently served by fiber[3]. With EoC, a huge incremental broadband capacity opportunity is opened up without huge fiber deployment costs. And as surveys have shown that Ethernet networks are at least 23% less costly to operate than ATM or TDM networks, the simplicity and consistency of Ethernet-over-Copper over 2BASE-TL bring large added savings in operating expenses to complement the capital savings and new broadband capabilities.

[3] According to Vertical Systems Group (www.verticalsystems.com).

What is Driving the Adoption of Ethernet-over-Copper?

Ethernet-over-Copper is deployed by multiple carriers – both incumbents and competitive carriers - in all major markets in the US and there are plans to roll out services to additional markets in 2009 and 2010. The adoption of Ethernet-over-Copper is driven by a confluence of enterprise needs: the near universal adoption of the Internet Protocol (IP), the economics of Ethernet technology deployment, and the vast availability of the copper asset.

Enterprises worldwide are requesting Ethernet based access services from their respective telecommunications providers. This demand is created as a result of several market forces, most notably:

- **Virtually all enterprise backbones are built using Ethernet technology.** IT professionals are very familiar with Ethernet technology, and their test equipment and training is geared towards Ethernet. During the past decade these professionals have realized significant cost savings, gains in throughput, and increases in reliability as a result of their Ethernet implementations and continued improvements in Ethernet technology. These Enterprise customers want to extend the same technology from their LAN to their WAN – simplifying and optimizing not only the LAN but also the entire extended enterprise.

- **Business customers continue to adopt IP-based applications**. These applications dramatically expand their consumption of packet network capacity. Existing business applications, file servers, intranet and training applications, email, digitized document management systems, and corporate Internet sites, continue to become more multi-media oriented. Corporate applications such as Storage Area Networks (SANs), offsite data replication, corporate video broadcast and video conferencing (which have traditionally used TDM circuits, B-ISDN, or dedicated fiber) are all moving toward the use of IP as their primary protocol. Lastly, large and small businesses are accelerating their adoption of Voice over IP (VoIP), with its economy, flexibility and call routing capability, to connect large offices, small offices, and even employees' homes. To deliver the quality of service expected by end users and demanded by packet data, voice and video applications, these services must be supported by end-to-end transport networks with high capacity, high availability, and high performance denoted by low packet loss, packet latency, and packet jitter.

- **Business customers look to Ethernet-based services for its cost advantages.** Although pricing varies from market to market, in many markets today, business customers pay ~$500 per month for a T1 based service operating at 1.54 Mbps. Up until now, when these customers wanted more bandwidth they had to purchase additional T1s in $500 increments per 1.54 Mbps T1. With the advent of Ethernet-based services and specifically Mid-Band Ethernet services, these same business customers can now order a 10 Mbps Ethernet service at less than $1,500 per month. With this service the business receives over 2X the bandwidth for the same price, and, most importantly, zero additional construction costs, and, the telecom carrier is able to generate more income per month for each customer. These services enable each customer to solve their bandwidth needs while reducing their overall cost per megabit for telecommunications services.

Carriers want to utilize Ethernet technology to meet the requirements of their customer base for more bandwidth and new services, and they want to reduce their operational expenses. Ethernet-based services provide "stickier" services for Carrier Service Providers. In today's market business customers are demanding more bandwidth and more services at a lower cost per megabit. Telecom carriers can only meet these demands with an established and ubiquitous technology like Ethernet. Furthermore, Ethernet services create a "stickier" customer experience. This stickiness primarily results from Ethernet's flexibility and customization options compared to legacy TDM-based approaches. With Ethernet carriers can cost effectively:

- Layer multiple services such as VoIP, private line, and Internet Access on one physical connection.
- Deliver bandwidth in 1Mbps increments and change levels of bandwidth delivery without rolling a truck to a customer location.
- Significantly reduce their overall operational expenses by as much as 80%. Ethernet is inherently a plug-and-play technology that requires much less provisioning and troubleshooting work than legacy TDM-based technologies.

Ethernet has long ago won the battle to become the natural link layer protocol for IP-based applications and services. IEEE 802.3 Ethernet standards have evolved to extend electrical interface speeds from 10 Mbps to 100 Mbps to 1000 Mbps, optical interfaces exceeding 10 Gbps, and wireless WiFi interfaces up to 52 Mbps. Interface cards and Ethernet switches are ubiquitous and offer very high capacity at a very inexpensive cost per bit handled, resulting in the near total domination of enterprise building LANs and campus area networks by Ethernet. Enterprises now wish to interconnect multiple sites and connect to the public Internet while maintaining the performance of their applications.

This market demand is being met by the implementation of high capacity, metro area and/or intercity Ethernet services at attractive price points by all major telecommunications Service Providers. In urban areas, core networks have been rapidly built out, and intensive capital spending programs have deployed fiber access to large buildings and major data centers. However, even in urban areas, nearly 80% of the business locations are not currently served by fiber[4]. Furthermore, the business case to invest in the capital cost to deploy fiber in the near term for one or two Ethernet service terminations is marginal at best.

Ethernet-over-Copper is ideal for today's economic climate. It leverages existing copper. It allows for "pay as you grow" deployment of electronics rather than up-front major capital and construction projects. And, as part of transparent Ethernet services, it can be operated in conjunction with the existing Ethernet management infrastructure. The footprint of the carrier metro and wide area Ethernet networks is extended, offering a high-margin revenue opportunity to the Service Provider.

[4] According to Vertical Systems Group (www.verticalsystems.com).

Target Applications

Business Access

The predominant initial application for Ethernet-over-Copper is the extension of the edge of the Service Provider's metro and wide area Ethernet networks to business customer locations that are not served by fiber. Businesses, government agencies, and other community anchor institutions, such as hospitals, schools, and libraries, are demanding higher bandwidth Ethernet-based services. These services are critical whether in the city center or in rural areas.

The Metro Ethernet Forum (MEF) has developed a broad set of Carrier Ethernet service definitions including the User Network Interface (UNI) point and point-to-point and multipoint connectivity between these interfaces. Ethernet-over-Copper can provide network edge access for:

- Private E-Line Service – A single point-to-point Ethernet virtual connection links two UNIs to deliver transparent LAN interconnection;
- Virtual Private E-Line Service – Multiple Ethernet virtual connections link two UNIs to provide connectivity for multiple VLANs;
- E-LAN service – Based on a multipoint-to-multi-point Ethernet virtual connection that bridges traffic between LANs in multiple offices;
- Symmetric commercial Internet access services.

These services may be delivered individually over an EOC access or, using the Virtual LAN capability of Ethernet, multiple services delivered over a single Ethernet-over-Copper edge connection. Furthermore, because these service definitions allow the access links to operate at different bandwidths, EoC may be provisioned with different data rates (due to differing loop length and condition) at each end, or may be used at lower speed to reach a number of branch offices while the main enterprise data center is served at a higher speed by fiber.

Mobile Wireless Backhaul

Mobile Wireless Service Providers' GPRS/EDGE networks and the evolution toward 3G wireless data services using CDMA2000 and EV-DO (Evolution Data Optimized) broadband technologies are changing the requirements for backhaul from the cell sites to the core network. A very limited proportion of cell sites are being served by fiber (less than 5%), most with only 1-2 T1/E1 backhaul links. Additional capacity is required for these growing data services, but TDM access does not meet the exponential bandwidth growth and resiliency needs of today's mobile wireless backhaul.

Ethernet-over-Copper can cost effectively deliver bandwidth to a cell site. Two pairs from the existing copper infrastructure are required for each T1/E1 to deliver 1.5Mbps/2Mbps respectively. Ethernet-over-Copper delivers 4 to 7 times more capacity over the same number of copper pairs. Increased capacity can be provisioned simply through the addition of bonded channels, growing up to 20 Mbps or more of low overhead Ethernet capacity. Unlike T1/E1s, the resiliency of the EOC bonded pairs infrastructure allows service to continue uninterrupted in the event of the failure of a pair whereas an NxT1/E1 solution results in an outage.

WiFi and WiMAX Backhaul

With the burgeoning deployment of public and private wireless data infrastructure, WiFi at 52Mbps with 802.11g and WiMAX on the horizon, comes an increasing need for backhaul of the data from base stations to the network core. Fiber deployment to an individual base station is not warranted by the bandwidth requirements, but the performance advantages of broadband wireless access are lost if a low speed T1/E1 connection back to the network is used.

Ethernet-over-Copper provides a bandwidth and cost appropriate solution to broadband wireless data backhaul. Rapid deployment of an initial backhaul link over one or two pairs is possible, with subsequent scaling up of the bandwidth being available through the addition of pairs to the bonded group as the usage from a particular base station grows.

DSLAM Backhaul

Residential Triple Play services (IP-TV, VoIP, and high speed Internet Access) are driving carriers to deploy IP DSLAMs further and further into the access network. While optical fiber deep into the access network is the ultimate goal (fiber to the neighborhood, fiber to the curb, fiber to the home "FTTx"), Service Providers need to be able to reach the greatest number of customers with broadband IP services as quickly as possible.

Ethernet-over-Copper is a natural complement to FTTx deployments, delivering native broadband Ethernet links to IP DSLAMs without costly, long lead-time fiber builds. Bonding pairs together to provide data rates up to 200Mbps (bonding up to 40 pair) over a resilient link provides broader service area coverage, sooner, resulting in shorter time to revenue for broadband services.

Ethernet-over-Copper Value Proposition in Brief

The value proposition for Ethernet-over-Copper is a combination of time to market, services and cost. Key aspects of the value proposition include:

- Cost and Reach
  - Copper is valuable nationwide ubiquitous infrastructure
  - Ethernet LAN technology has achieved low cost through commoditization
  - Available to over 95% of un/underserved community anchor institutions
  - Price – better price per bit than the alternatives
- Services
  - Ease of Use - Common technology with the LAN, no special hardware or operational knowledge required
  - Breadth of Services - Offer point-to-point or multipoint services
  - Capacity available in the 2-20 Mbps range
  - Reliability – natural redundancy/resiliency of the pair bonding scheme
  - Service Level Agreements – for performance and availability
- Time to Market
  - Short provisioning cycles (no network build required)
  - Growth ready - speed upgrades with no need to turn down service

Several market analysis organizations such as Vertical Systems Group (www.verticalsystems.com), the Yankee Group (www.yankeegroup.com), Heavy Reading (www.heavyreadying .com), and Infonetics (www.infonetics.com) have identified the large and rapidly expanding Ethernet Service Market. Based on Vertical Systems estimates that only 20% of business locations have access to fiber, Ethernet-over-Copper will be a large part of that market.

Conclusion

Telecommunications traffic, for voice, video, and data applications is moving to a fully packetized world. Ethernet long ago won the battle for technical supremacy as the protocol for delivering these packets. Large, medium and small Enterprise customers, government agencies, non-governmental organizations, and wireless carriers are all looking for economic, high capacity metro and wide area alternatives to tie their locations together. They have found T1 speeds insufficient, yet 100 Mbps to 1 Gbps fiber services are not affordable for the majority of locations.

Ethernet-over-Copper is an enabling technology that fills a huge gap in the ability to provide broadband services to anchor tenants in un/underserved areas – by most estimates over 80% of buildings in the US are not reachable by fiber[5]. It enables broadband services to unserved and underserved areas by: 1) directly delivering business-grade Ethernet access to community anchor institutions, 2) providing high speed connectivity to rural cell towers, and 3) providing a wider reach for residential DSL services.

EoC delivers business access at rates between 2Mbps and 50 Mbps at a cost significantly below other access services. EoC offers higher ROI (reaching anchor tenants at 1/10th the cost of fiber), can be rapidly and cost effectively deployed over the existing copper infrastructure (10x faster than pulling fiber), and, as native Ethernet technology, can be readily integrated into anchor tenant networks and carrier Ethernets. EoC allows the provider to offer customers a consistent, universal Ethernet offering to any location.

[5] According to Vertical Systems Group (www.verticalsystems.com).

Appendix A - Ethernet-over-Copper Technology Handbook

IEEE 802.3ah 2BASE-TL
for Executives

MID-BAND ETHERNET TECHNOLOGY HANDBOOK

The New Standard for
High-Speed Symmetrical
Copper Access

HATTERAS™
NETWORKS

Foreword

There is no doubt that Ethernet over copper is becoming a significant contributor to telecom carriers' bottom line, as it is enabling them to cost-effectively mine their copper infrastructure for new customers and enhanced service offerings.

The ratification of the IEEE 802.3ah Ethernet in the First Mile (EFM) standards was a watershed event for the industry, prompting leading carriers worldwide to move from passive observation to active participation in RFPs, trials, and deployments. The year 2006 is going to be the turning point in how carriers generate additional revenue from their copper networks in advance of, and many times in lieu of, fiber roll outs. The decisions carriers make today on how they build their last mile service infrastructure could easily dictate their ability to compete and survive as the competition for broadband business customers heats up.

Hatteras Networks played a key role within the IEEE 802.3ah Task Force, contributing ideas, market knowledge, and technical direction as these market changing standards were being set. The following "The Mid-Band Ethernet Technology Handbook: IEEE 802.ah EFM 2BASE-TL for Executives," is a must read for telecom decision makers that need to understand not only the technology, but the market potential for Ethernet over copper solutions, and why so many carriers are already deploying the Hatteras solution and improving their bottom-line.

- Michael Howard,
Principal Analyst & Co-Founder, Infonetics

Contents

Executive Summary

In June 2004, IEEE 802.3 ratified a new amendment to the Ethernet standard – IEEE 802.3ah Ethernet in the First Mile (EFM). This standard adapted Ethernet – the best known and most widely used LAN technology in history – for widespread deployments in carrier access networks. With EFM, complex and costly ATM or SONET/SDH access networks can be migrated to simpler, more cost-effective Ethernet access networks, resulting in immediate savings in capital and operating expenditures, as well as increased bandwidth and service options to the subscriber.

As part of its sweeping potential in the access network, the EFM standards group defined two technologies for delivering Ethernet over plain old telephone lines: 2BASE-TL and 10PASS-TS. These technologies offer higher bandwidth and higher quality services than existing T1/E1 and xDSL solutions, delivering the simplicity and flexibility of Ethernet, while still maintaining full spectral compatibility within the existing network. And since these standards were created, their capabilities and benefits have been accepted by carriers and standards organizations across the globe.

2BASE-TL and 10PASS-TS are revolutionizing and expanding the copper access network. For those carriers offering Ethernet services over optical or SONET/SDH infrastructures, these Mid-Band Ethernet™ technologies make Ethernet services available to the vast majority of customers that do not have access to fiber. Instead of Ethernet services being limited by fiber availability to less than 10% of potential business sites, the services are now available to almost any subscriber location. With distance potential beyond 20 Kft (6 Km), 2BASE-TL can reach almost any business subscriber, providing a universal multi-megabit on-ramp to your Metro Ethernet network.

For those carriers already delivering services via ATM-based xDSL technology, Mid-Band Ethernet provides a path to simpler networks with lower operating expenses, and to differentiated, higher-margin services currently out-of-reach using existing technologies. The simplicity and cost-effectiveness of Ethernet yields immediate savings in capital and operating expenditures. And with purpose-built Mid-Band Ethernet platforms, VoIP and VPN services can be delivered right alongside simple high-speed Internet access, yielding additional layered revenue from the same customers, all while maintaining a consistent access network and service set for all customers. Subscribers connected via 1000BASE-X Gigabit Ethernet and 2BASE-TL experience the same service and are managed with the same paradigms and the same software – it's all Ethernet. The only difference is the access media and the available symmetrical bandwidth.

Adhering to the EFM standards is just the first requirement for providing a Mid-Band Ethernet solution. Additional capabilities are needed to reach the full potential of 2BASE-TL services, including full service edge functionality, with QoS, carrier grade reliability, and a low initial cost.

The next few pages dive into the details of the EFM standards, showing how they work, and how they can help you increase revenues, decrease costs, and move toward a simplified network infrastructure.

Background and Introduction

A dominant force in the LAN for decades, in recent years Ethernet has exploded into the provider network as the preferred layer-two transport for core and metro networks. Perfectly built for next-generation IP and VPN services, Ethernet yields simplicity for easy deployment, flexibility for a variety of services, and lower cost for large-scale networks.

The Ethernet in the First Mile Project (IEEE 802.3ah)

The demand for Ethernet in the carrier network was realized within the IEEE 802.3 Ethernet standards body, which created the 802.3ah Ethernet in the First Mile (EFM) task force. This task force developed enhancements to existing Ethernet standards targeted to address carrier concerns about deploying Ethernet in the access network. EFM quickly focused on four separate problem areas that required new standards development

1. New optical Ethernet physical layers with the reach, bandwidth, and environmental characteristics for outside plant deployments.
2. An Ethernet passive optical network (E-PON) technology for residential deployment of triple-play (voice, video and data) services.
3. Operations, administration, and maintenance (OAM) functions to simplify monitoring and troubleshooting geographically dispersed Ethernet networks.
4. New physical layers for Ethernet over standard telephony wiring for universal broadband Ethernet coverage.

Within the fourth category of EFM work items, the EFM task force developed two new Ethernet physical layers to address two very different and very important segments of the market. The first technology, 2BASE-TL, is a long-reach Ethernet-over-copper technology focused on high-bandwidth symmetric services for business customers from central offices or remote terminals. 2BASE-TL is the natural upgrade and replacement for today's T1/E1 and G.shdsl services. The second technology, 10PASS-TS, is a short-reach high-bandwidth asymmetric technology targeted for in-building or FTTC deployments, permitting some increased bandwidth options for residential or business access. Both 2BASE-TL and 10PASS-TS enable native Ethernet frames to move across existing voice-grade copper pairs in carrier access networks.

What's in a Name?

The IEEE 802.3ah copper pair standards are each named in the traditional Ethernet fashion with 2BASE-TL and 10PASS-TS. As with all Ethernet interfaces, the naming of the technology tells you a lot about the technology itself. Let's dissect one of these names for its true meaning.

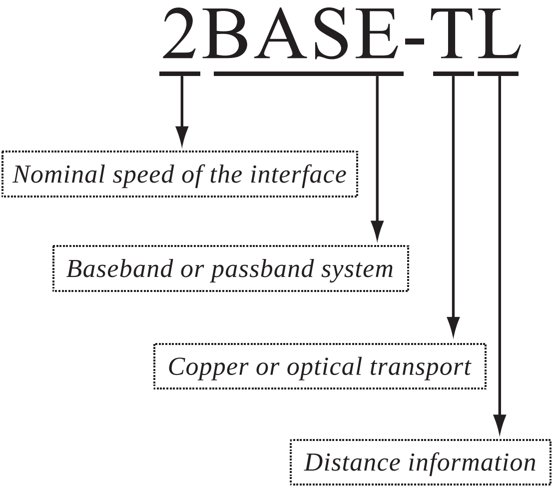

The first digit gives the speed of the interface. In the case of 2BASE-TL and 10PASS-TS, this is considered the nominal speed of the interface as the interface itself can run at different speeds depending on the conditions of the environment. 2BASE-TL is targeted to deliver at least 2 Mbps symmetric operation per pair in a long-reach nominal noise environment, while 10PASS-TS is targeted to deliver at least 10 Mbps per pair in a very short-reach nominal noise environment. The second token is BASE or PASS, indicating whether the system is baseband or passband. 10PASS-TS is a passband system in that it can run on copper pairs that have a baseband POTS service (e.g., traditional POTS voice with data on the same line). 2BASE-TL is baseband in that it cannot run on the same line as POTS, but can carry VoIP. Both systems are run over twisted copper pair (hence the T). Finally, 2BASE-TL is a long reach service, while 10PASS-TS is a short reach service (hence the L and S).

And there you have the parsing of the IEEE 802.3ah nomenclature!

The EFM Family

The EFM technologies cover the full spectra of copper access deployment possibilities, from short-reach to long, from business to residential. When looking at rate/reach possibilities, 2BASE-TL and 10PASS-TS combine to blanket the chart as shown in Figure 1.

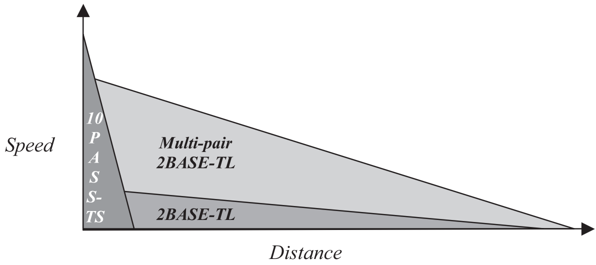

Figure 1. EFM Family Rate/Reach Applicability

With 10PASS-TS, EFM can offer very high rates on very short loops, reaching as high as 100 Mbps in asymmetric mode, and 50 Mbps in a more symmetric mode. 10PASS-TS also supports a bonded multi-pair application for additional bandwidth and reliability. It is targeted at 5 Kft (1500 m) and below in terms of Customer Serving Area (CSA) distances, so the reach is very limited. Additionally, although 10PASS-TS has more symmetric modes of operation, the reach of the symmetric capabilities is even more limited than the asymmetric reach. For these reasons, 10PASS-TS targets true triple play residential services, with asymmetric, very high-bandwidth service potential.

In contrast to 10PASS-TS, 2BASE-TL is focused on delivering symmetric services to business customers from central offices and remote terminals. It supports service delivery out to CSA distances (9-12 Kft, 2700-3600 m) and beyond. With a maximum symmetric rate of 5.7 Mbps per pair, 2BASE-TL delivers high-bandwidth Ethernet services over just a single pair. However, EFM also defined a new multi-pair bonding technique whereby up to 32-pair can be grouped into a single Ethernet port, thus increasing the bandwidth and resiliency of the subscriber connection. Businesses can now be reached with symmetric services of 10 Mbps and higher, on as few as two pairs of copper, as 2BASE-TL continues to be rolled out as the next generation replacement for traditional T1/E1 services.

The EFM standard leveraged the latest DSL layers as defined by the International Telecommunications Union (ITU) as the physical layers for Mid-Band Ethernet. By utilizing these existing standards, IEEE 802.3ah continues to benefit from the high volume of DSL chipsets, while significantly improving upon the original silicon by defining new and efficient mechanisms for Ethernet transport.

History of DSL

Two standards bodies have long studied how to utilize telephony copper wiring for data transport. From early analog modems thru ISDN, including all of today's advanced modulation techniques, the ITU and ANSI T1 have driven international standards for copper access networks. Figure 2 shows a historical timeline of ITU/T1 advancements in digital access over twisted pair.

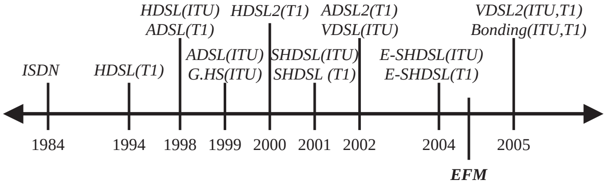

Figure 2. Historical DSL Developments

2BASE-TL

As noted in a previous section, 2BASE-TL offers a nominal symmetric bandwidth of at least 2 Mbps in a reasonable environment at reasonable distances. 2BASE-TL is based on the same physical layer as the Enhanced SHDSL standards of ITU and ANSI T1 (also known as G.991.2.bis or E-SHDSL). Whereas SHDSL (Symmetric High-speed DSL, G.991.2) had a maximum symmetric rate of 2.3 Mbps, Enhanced SHDSL can run up to 5.7 Mbps on a single pair. With such high-speed symmetric access, subscribers can be offered a true 10 Mbps Ethernet service on as little as 2-pair.

2BASE-TL and Enhanced SHDSL increased the bandwidth potential of SHDSL in two dimensions. First, a second constellation (or symbol encoding) is allowed which increases the throughput by 33% without affecting the spectral properties of SHDSL. This additional higher constellation cannot be used on the longest loops, but does provide a spectrally "free" throughput increase on loops up to 10 Kft (3 Km), depending on the noise environment. Second, 2BASE-TL and Enhanced

SHDSL increase the frequency (number of symbols per second) as compared to SHDSL, thus allowing even more throughput. This frequency addition increases the noise created by the technology, but it still falls within North American and most international spectral guidelines such as ANSI T1.417.

10PASS-TS

The EFM short-reach solution is based on VDSL (Very high-speed DSL). One of the major technical decisions of the EFM task force was to decide which VDSL technology was best suited for the short-reach Ethernet physical layer. At the time, there were two VDSL candidates. One candidate was based on Discrete Multi-Tone modulation (DMT), and the other was based on Quadrature Amplitude Modulation (QAM). Both technologies could yield similar performance results yet only one could be selected. Until EFM forced a decision, both technologies had progressed equally through ITU and ANSI T1 standards bodies, with no organization able to select a technically superior solution.

After many months of debate, the EFM task force voted to use VDSL-DMT as the physical layer for 10PASS-TS instead of VDSL-QAM. The hope that VDSL-DMT could leverage the technology and volume of ADSL (which is also based on DMT technology) was a key factor in the selection process.

Carrying the Torch Beyond IEEE 802.3ah

In the time since IEEE 802.3ah developed the advanced mechanisms for Ethernet transport and bonding, many other standards bodies around the world have recognized their work by incorporating those same techniques into other international standards.

Both ANSI T1 and the ITU have referenced the IEEE 802.3ah techniques for all forward looking DSL technologies. The highly efficient IEEE 802.3ah method for framing and transporting Ethernet packets on xDSL lines has been incorporated into ADSL2 and VDSL2 as the preferred packet transport technique. The dynamic and flexible IEEE 802.3ah methods for bonding multiple pairs has been standardized by both groups as the method for delivering packet transport over more than one copper loop.

Not only have these groups standardized on the IEEE 802.3ah methods, they have also worked to improve those methods. For example, the IEEE 802.3ah framing mechanism has been extended by the ITU to allow transmission of small (less than 64-byte) frames. This simple adaptation of the IEEE method now allows for the use of the same technology for native IP transport (where frames may be very small).

The fast and widespread technical and market adoption of IEEE 802.3 ah, as well as the dedication to improving the technology, has cemented IEEE 802.3ah as the best way to deliver Ethernet services over the copper loop infrastructure.

Spectral Compatibility and International Application

One of the key issues with any technology operating over telephony wiring is the effect of the technology on other services. Telephony cables are deployed in binder groups of tens of pairs bound together in a common outer sheathing. When a service is deployed on a pair, it creates noise that affects other pairs in the binder group, and sometimes in nearby binder groups. The impact of a technology on other services in a binder group is measured by its spectral compatibility.

There is no hard and fast rule as to what is "too much" noise. In different parts of the world, different kinds of cables are deployed, the loops are of different lengths, and the deployment practices are different. As such, what is acceptable in one area may be completely unacceptable in another. Each country is therefore free to define its own spectral compatibility guidelines for services deployed in its telecommunications infrastructure. In North America, for example, the spectral deployment guideline specification is known as T1.417.

As an international standard, it is important for Ethernet to be deployable anywhere in the world. Each of the EFM technologies (2BASE-TL and 10PASS-TS) are based on basis systems, which are technologies that are universally deployable throughout the world. These technologies are capable of operating under different spectral guidelines depending on where in the world they are deployed. Different spectral guidelines yield different performance results, so the effective throughput of the technology is limited by the governing spectrum rules of the local country. EFM technologies are internationally deployable anywhere in the world provided they are configured to conform to the regional spectrum guidelines.

Carrying Ethernet Packets Over Copper

A long-standing tradition in Ethernet is that the method for carrying the actual frames over the wire must be

(a) low overhead, and (b) incredibly resilient to false packet acceptance. False packet acceptance (FPA) is the probability that when a corrupted frame is received, the corruption is not detected. As an example, 10BASE-T uses what is known as 4b/5b encoding, where every 4-bits of data is encoded as 5-bits on the wire for additional resiliency, with 20% overhead (1 of every 5 bits is overhead). Later Ethernet technologies like 10G Ethernet use 64b/65b encoding (1 out of every 65 bits is overhead), resulting in improved efficiency as more and more speed is squeezed out of the medium.

The Mid-Band Ethernet technologies use a novel encoding scheme called 64/65-octet encoding – where there is one overhead byte for every 64 bytes of data. This encoding scheme is incredibly efficient, which is vital in access technologies that must adapt to the environment to deliver the highest possible speed given existing outside plant conditions.

Additionally, 64/65-octet encapsulation includes measures to improve the false packet acceptance results of traditional DSL encoding. DSL physical layers generally operate in modes that yield a bit-error rate of 10-7. Traditionally, Ethernet technologies (and the IP layers above them) have been built upon an architecture where false packet acceptance cannot probabilistically occur. To achieve FPA performance acceptable for Ethernet and IP delivery, the 64/65-octet layer appends every frame (or fragment) with a CRC in addition to the Ethernet FCS. The combination of these two error-checking codes practically eliminates the possibility of FPA, thus maintaining the historical reliability of Ethernet.

These benefits result in a more efficient and a more reliable access network technology. As an example, carrying Ethernet over ATM results in 20-50% overhead, while carrying Ethernet natively via 2BASE-TL and 10PASS-TS results in less than 5% overhead. This allows carriers to squeeze more bandwidth (and more revenue) out of their existing infrastructure. The days of throwing bandwidth away on ATM or TDM overhead are officially over.

Mid-Band Ethernet: Multi-Pair Aggregation

EFM copper technologies introduced a novel mechanism for utilizing multiple copper pairs to deliver additional bandwidth to the subscriber. The loop aggregation techniques developed in 802.3ah are able to optimize the utilization of the set of lines as well as make adding and removing pairs to the aggregate a snap!

How It Works

IEEE 802.3ah developed a loop aggregation technique for Ethernet optimized for copper access. As shown in Figure 3, the aggregation technique is transparent to higher layer applications as it sits below the Ethernet MAC. The switching and services layer of the device can be consistent across optical, CAT5, and EFM Ethernet interfaces, allowing the provider the ability to offer a consistent service offering over any type of access media.

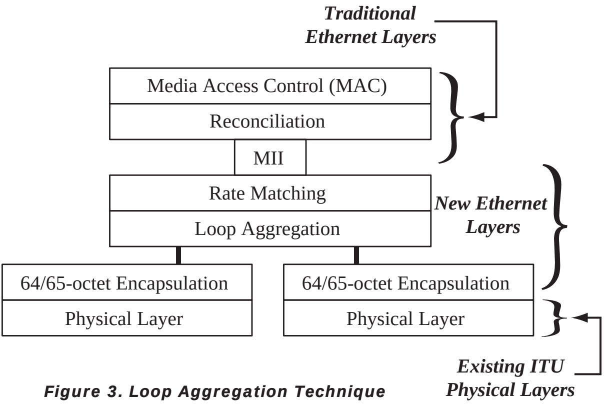

Figure 3. Loop Aggregation Technique

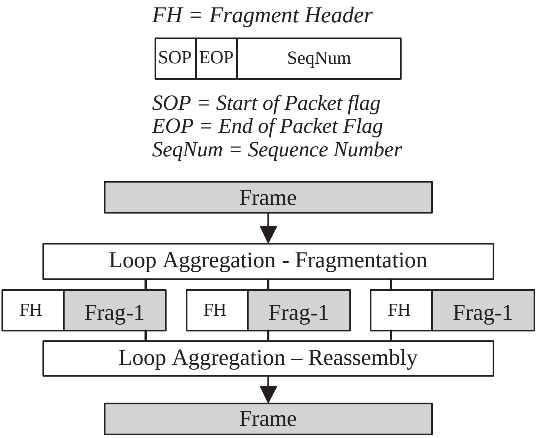

Figure 4. Fragmentation Header is Prepended

The loop aggregation techniques of IEE 802.3ah are both simple and powerful. Frames are passed to the loop aggregation layer from the higher layer, where they are fragmented and distributed across the loops within the aggregate. When transmitted across the individual loops, a fragmentation header is prepended (see Figure 4), which includes a sequence number and frame markers. This header is used by the receiver to resequence the fragments and to re-assemble them into complete frames. To allow vendor differentiation, the algorithm for partitioning the frames over the loops is not specified. However, it must obey certain rules in that fragments must obey size constraints, and the loops in an aggregate must obey rate and differential delay constraints. As long as the loop aggregation algorithms obey these constraints and restrictions, any fragmentation algorithm can be handled by the reassembly process, yielding a very flexible and interoperable solution.

Not Link Aggregation

Although they may look similar, Loop Aggregation as defined in 802.3ah is very different than Link Aggregation as defined in 802.3ad. Loop aggregation fragments individual frames into variable-sized segments in order to minimize latency and maximize utilization of disparate speed links. Link aggregation load-balances frames over equal speed links in order to increase aggregate throughput.

One very key difference is that the loops with Loop Aggregation (802.3ah) can be running at very different speeds – something not possible with 802.3ad Link Aggregation. Likewise, the ability to fragment large frames into smaller pieces is very important when trying to minimize latency. A 1500 byte Ethernet frame takes 12ms to transmit when lines are running at 1 Mbps. Breaking this frame up into N equal size fragments decreases transmit latency for this frame by a factor of N.

Overcoming IMA Bonding Limitations

Inverse Multiplexing over ATM (IMA) is another traditional technique used to aggregate multiple lines together. IEEE 802.3ah Loop Aggregation improves upon IMA in several key aspects. First, IMA is limited to links of the same speed as it uses a round-robin distribution algorithm. In realistic copper environments, loops train at different speeds, sometimes as much as 3 or 4 times the rate of one another, even when going between the same two endpoints. This is because the noise on each pair can vary greatly. An IMA solution "dumbs down" the faster pairs to the slowest speed, thus wasting a large portion of the potential bandwidth.

Additionally, IMA is ATM, not Ethernet, and carries with it the overhead and complexity of ATM cells and virtual circuits. With the minimal Ethernet overhead, the perceived user thoughput can be 200% of the higher overhead ATM solution.

Automatic Resiliency for the Most Demanding Subscribers and Applications

In addition to the efficiency and performance benefits of 802.3ah, loop aggregation has the added benefit that it's automatic. Pairs can come and go, and the Ethernet interface remains operational — only the available bandwidth is affected. New pairs can be wired up and automatically joined to the aggregate group with no additional configuration, realizing the plug-and-play potential of Ethernet.

The resiliency of 802.3ah loop aggregation can satisfy the most demanding business customers and support any application. When a pair fails, that pair is detected and removed from the aggregate in just a few milliseconds. Established Voice-over-IP calls remain operational, and the callers don't even notice that a problem has occurred. Video streams continue to play as if nothing changed. Applications, and their users, will not be able to detect that one of the pairs has failed except by the loss of some bandwidth. And when that pair comes back online, it is seamlessly added to the aggregate, and that again goes unnoticed by the applications and users. This makes IEEE 802.3ah the most suitable technology for the business services of today and tomorrow, where unreliable, best effort delivery is simply not enough.

Hatteras Networks' Ethernet Service Edge Solution offers much more than "just another DSL." It boasts a service-enabling architecture that yields a highly efficient network capable of deploying the services of today and tomorrow, all while simplifying deployment with new plug-and-play features.

Not Just Another DSL

The Ethernet standards leverage well-known and well-tested physical layer standards of existing DSL technologies in order to minimize cost and complexity, maximize interoperability and simplicity, and satisfy spectral compatibility. However, they are not simply another xDSL - any more than optical Gigabit Ethernet is simply another SONET/SDH because they use the same optical wavelengths and fiber specifications.

Per-pair Efficiency

Delivering native Ethernet over existing copper plant brings simplicity and flexibility to deployments. Encapsulating Ethernet over DSL has traditionally been an extremely complex process with very high overhead as shown in Figure 5. In traditional deployments, Ethernet has been encapsulated in the Point-to-Point Protocol (PPP), then segmented/re-assembled with ATM AAL5, run through a fast/slow path DSL multiplexing layer, and finally over the DSL modem where data is transmitted over the wire with all of the added overhead.

The overhead from ATM alone results in 20-50% inef-

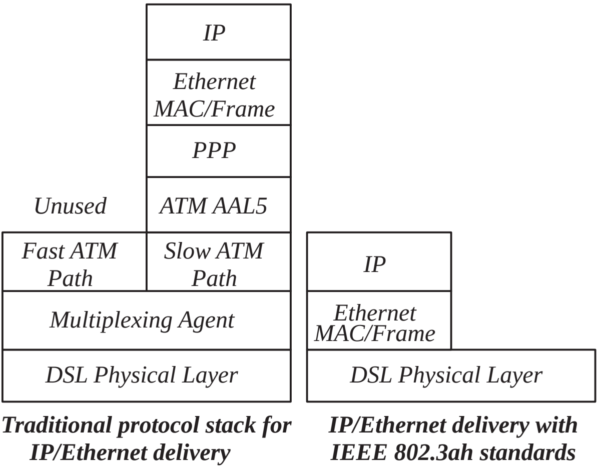

Figure 5. Delivering IP/Ethernet: Old Versus New

ficiency when carrying Ethernet traffic – this is additionally compounded with additional encapsulations most often used (such as PPP). As an example, consider the case where 4 physical lines train at 1 Mbps, 2 Mbps, 3 Mbps, and 4 Mbps, respectively. This yields an aggregate physical bit rate of 10 Mbps. When running Ethernet over ATM on these lines, the ATM overhead alone eats up 20-50% of the available bandwidth – the carrier cannot guarantee more than a 5 Mbps service and the service will on average yield less than 8 Mbps. However with Mid-Band Ethernet, the efficiencies of 2BASE-TL and 10PASS-TS allow the carrier to guarantee a much faster service to the user, averaging at least 9.7 Mbps to the user. This increase in bandwidth is a significant improvement in the bandwidth bottleneck of the access network.

Bonding Efficiency

Continuing with the previous example, an ATM IMA solution would effectively limit each of the 4-pairs to the lower speed, which in this case is 1 Mbps. Thus the 10 Mbps would be reduced to 4 Mbps, before the 20-50% overhead is applied. This would effectively result in a service of between 2 and 3.2 Mbps to the user. However, with the ability of the Mid-Band Ethernet technologies to utilize such disparate speed links, the resulting service would yield closer to 9.7 Mbps to the user. The differences are dramatic: 3.2 Mbps versus 9.7 Mbps. And that doesn't even account for the overhead of the PPP connection in the ATM architecture.

Simplicity and Consistency

Perhaps the most significant benefit of Mid-Band Ethernet is the simplicity it brings, and the resultant savings in capital and operating expenses. Surveys have shown that Ethernet networks are at least 23% less costly to operate than ATM or TDM networks. As true Ethernet technologies, 2BASE-TL and 10PASS-TS can yield immediate large operational savings.

Additionally, 2BASE-TL and 10PASS-TS yield consistent next-generation architectures across any medium. For optical customers, the generally preferred access mechanism is gigabit Ethernet. For customers in a high-rise, 100BASE-T is a popular interface. As operators design and deploy services over such technologies, it is ever important that these same services are available on the copper access network. Using 2BASE-TL and 10PASS-TS, the services and management can be the same – it's an all Ethernet network. The only difference between customers connected via fiber and customers connected via twisted pair is the potential bandwidth available to the customers. The services and management remain consistent between headquarters and branch office locations of the customer, allowing the provider to capture the whole customer with a differentiated, high margin service set.

There were two key decisions made in IEEE 802.3ah related to which DSL technologies to use as the underlying physical layers. One decision was whether to use VDSL-QAM or VDSL-DMT as the basis for 10PASS-TS. As discussed earlier, technical and marketing reasons led the task force to select VDSL-DMT as the technology for short-reach Ethernet transport.

For the long-reach technology (2BASE-TL), 802.3ah had to decide between using Enhanced SHDSL (E-SHDSL) and ADSL2/plus. Although both technologies offer very good performance at the desired rates and reaches, traditional SHDSL has been successfully used for symmetric business services for several years, while ADSL has been deployed almost solely for residential applications. Each technology brings to the table its own advantages.

In the end, it was decided that 2BASE-TL should be based on E-SHDSL. ADSL2/plus, although a great fit for residential services, is not a natural fit for symmetrical business applications. The IEEE based this decision on the following reasons:

1) E-SHDSL is a naturally symmetric technology. As 2BASE-TL is targeted at next generation business access, symmetric performance was a defining criteria, and the symmetry of E-SHDSL was an advantage over the asymmetric nature of ADSL2/plus.

2) Many large and small operators backed the use of E-SHDSL as the best symmetric technology for this application. The supporting operators included RBOCs, PTTs, IXCs, and CLECs from across the globe.

3) E-SHDSL was designed for the business environment. The cable make-up for business services is very different than that for residential services. Business cables contain T1/E1 circuits, HDSL, HDSL2, and SHDSL disturbers. These pre-existing symmetric technologies create a noise environment in which ADSL2/plus does not offer good upstream performance. SHDSL and E-SHDSL, on the other hand, were designed for business environments and are much more robust against such disturbers.

4) Common business applications like radio signals (e.g., police radios) can cause interference with ADSL2/plus signals. This results in less reliable services, or in "spectral notches" around these interferers, which greatly decrease the bandwidth potential. SHDSL and E-SHDSL are immune to these interferers.

2BASE-TL Has Better Symmetric Performance in the Worst Case

To illustrate the performance of E-SHDSL and ADSL2/plus in the business environment, Figure 6 presents comparative symmetric performance results. For this figure, the potential symmetric bandwidth of ADSL2/plus and E-SHDSL was measured against a wide variety of noise conditions as specified in the appropriate standards. The noise conditions included everything except the common business T1/E1 disturber. For each distance, the worst-case symmetric throughput over all of the noise models was graphed. From Figure 6, you can see that E-SHDSL has better worst-case symmetric performance than ADSL2/plus. Note in some cases ADSL2/plus had better asymmetric performance, and in other cases ADSL2/plus had better best-case performance. But it never offered better worst case symmetric performance. Unlike residential best-effort Internet access where the best downstream burst speed is all that matters, for the business application of 2BASE-TL, it is important to deliver symmetric throughput in the worst case, so that business requirements and business SLAs can be specified and achieved under any conditions.

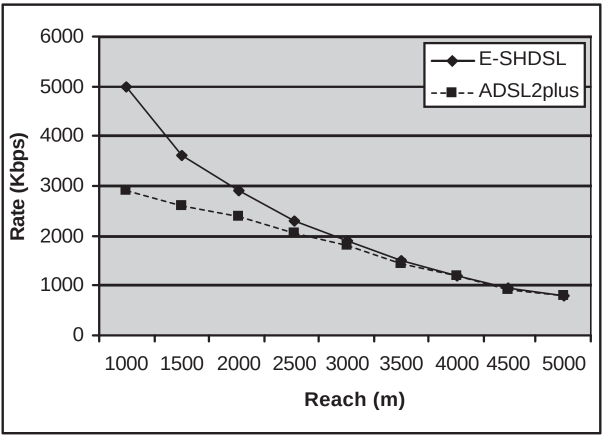

Fig 6. Worst Case Symmetric Performance <u>Without</u> T1 Disturbers

2BASE-TL Has MUCH Better Performance in the Business Environment

In addition to the superior worst-case performance against non-T1/E1 disturbers as shown in the preceding graph, E-SDHSL has significantly better resiliency against the most common business disturber – the T1/E1. Many businesses today use T1/E1 as the access

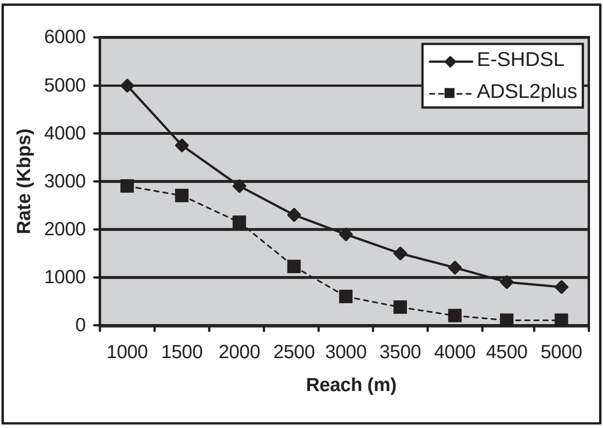

Fig 7. Worst Case Symmetric Performance With T1 Disturbers

method for their voice services, and for many of their data services as well. This disturber, which is not common in residential applications, has a significant impact on symmetric ADSL2/plus performance as shown in Figure 7.

So ADSL2/plus, though a very promising technology for the residential market, is not the best choice for symmetric business applications. Enhanced SHDSL proved a much more robust technology in the business environment, especially when T1/E1 disturbers are included in the comparison. As a next-generation T1/E1 replacement technology, the resilience, reliability, and symmetric performance in the business environment drove the selection of E-SHDSL as the foundation for 2BASE-TL.

VDSL2 did not exist when the IEEE 802.3ah was being developed. In fact, VDSL2 wasn't standardized until a year after the IEEE 802.3ah standard was ratified - long after the technical decisions were made. If VDSL2 were ready one year earlier, it would have been the underlying technology for 10PASS-TS.

Improving on 10PASS-TS

VDSL2 offers significant improvements over the original VDSL on which 10PASS-TS was based. In fact, because of the benefits of VDSL2 over VDSL and its timing in coming out shortly after IEEE 802.3ah was finished, true 10PASS-TS implementations are unlikely to ever be developed. However, as mentioned earlier, the techniques of IEEE 802.3ah copper access, including efficient transport and dynamic bonding protocols, have been carried over into other ITU and ANSI T1 standards. One can therefore use VDSL2, instead of VDSL, as the basis for a 10PASS-TS replacement, creating a new and improved short-reach Ethernet-over-copper technology.

VDSL2 Still Has Limited Symmetric Performance

Note, however, that a VDSL2-based 10PASS-TS implementation still suffers from the same limitations as a VDSL-based implementation – just not quite as much. For example, VDSL2 still has extremely limited reach, and even more so when deployed in a more symmetric manner. Figure 8 shows the upstream line rate performance of ADSL, VDSL2, and 2BASE-TL in a typical noise environment.

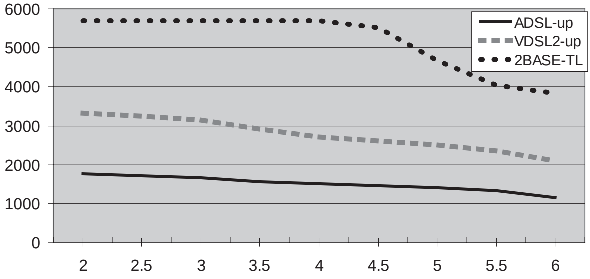

Fig 8. Worst Case Symmetric Performance With T1 Disturbers

Figure 8 assumes 26 AWG cable and 24 disturbers. The 2BASE-TL numbers assume a worst case 24-disturber model (all self disturbers), while the VDSL2 and ADSL numbers are not worst case (12 self and 12 HDSL disturbers).

As the figure shows, VDSL2 (and ADSL) upstream performance cannot support high-speed symmetric business applications with anywhere near the capabilities of 2BASE-TL. So although VDSL2 is a dramatic technology for triple-play residential services, it has a difficult time for business applications that have a hard requirement for symmetric capability.

VDSL2 Is Still Not a Fit for the Enterprise Market

The performance of VDSL2 suffers even more when classic T1/E1 disturbers are present in the environment, as was the case of ADSL2/plus as discussed earlier. As these classic business technologies are still very common to binders that serve enterprise locations, VDSL2 performance will commonly be much less than the diagram predicts.

For reaches longer than the 6 Kft shown in Figure 8, the relative performances of VDSL2 and 2BASE-TL are even more dramatically different. VDSL2 bandwidth drops to nearly zero very quickly, while 2BASE-TL provides a graceful decrease in bandwidth and can deliver services to 20 Kft and beyond. Therefore not only is the symmetric performance of 2BASE-TL superior under almost all conditions, the reach of the technology, and therefore the size of the addressable market, is significantly higher with 2BASE-TL than with VDSL2.

The size of the addressable market, and the importance of the ubiquity of the service, is something that cannot be underestimated. The distance distribution of subscribers from the central office differs from carrier to carrier, and from region to region. However, some tendencies remain constant:

- Enterprise customers are generally much closer to central offices than residential customers.
- The vast majority of enterprise customers (70-90%) are within CSA distance (12 Kft) of the central office.
- Less than half of enterprise customers are within the 5 Kft or so reach of VDSL2.
- There is some non-zero percentage of customers that remain at very long distances (18 Kft or beyond).

So although the exact distance distribution may differ depending on the geography, what is a constant is that 2BASE-TL can serve the vast majority of customers with a symmetric 10 Mbps service, while VDSL2 cannot reach even half of the customers with a much lower service offering.

Additionally, 2BASE-TL is a technology that can be regenerated in the outside plant to provide almost unlimited reach. Because 2BASE-TL is based on the well established standards of SHDSL and Enhanced SHDSL, and these standards are designed to allow regeneration in the outside plant, 2BASE-TL repeaters can be deployed to reach almost anywhere. VDSL2, as a very high speed short-reach technology, cannot offer the possibility of universal reach because it cannot be regenerated.

Note however that VDSL2 and 2BASE-TL are still very complementary technologies, just as was originally true with 10PASS-TS and 2BASE-TL. The two technologies target different reach segments, different market segments, and different symmetries. And when loops are very short (less than the 2 Kft lower limit showed in the diagram), VDSL2 can even deliver more bandwidth on a single pair. VDSL2 thus provides a more optimal technology for very short reach applications, and 2BASE-TL and VDSL2 can be used to offer complementary services across enterprise and residential applications.

A major market motivation for Mid-Band Ethernet is the ability to deliver Ethernet using the existing infrastructure – without adding fiber, without grooming the outside plant, and without redesigning the network. Carriers must reach the customers no matter where they are. Mid-Band Ethernet services based on 2BASE-TL can provide almost unlimited reach, similar to that of a traditional T1 or E1 service, so that Ethernet services can be delivered to any customer, anywhere.

Regenerating 2BASE-TL

All communications signals degrade over distance and time. The distance over which a signal can be transmitted and received without error provides a limit on the deployment capabilities of that technology. For example, optical Ethernet signals can go for hundreds of kilometers, while 10/100/1000BASE-T Ethernet can traverse just 100m. Mid-Band Ethernet sits between these two, with VDSL or VDSL2-based technologies having a reach of 1000-1500m, and 2BASE-TL reaching as far as 6km.

2BASE-TL is based on E-SHDSL, which has the important property that it can be easily regenerated in the outside plant. For years, line powered repeaters have existed for T1/E1 based services. These repeaters sit in the midst of the copper plant between the carrier and the customer, and regenerate the electrical signal to provide additional reach. This regeneration extends the reach of the technology by converting the line power to transmit a new signal, just as the signal is initially generated from a central office or cabinet.

Other xDSL technologies cannot be so easily regenerated in the outside plant. Technologies such as ADSL and VDSL2 do not have this capability – it is only the symmetric technologies, such as HDSL, HDSL2, and E-SHDSL that support the nearly unlimited reach offered by repeaters.

In practice, of course, the reach is not truly unlimited. Each repeater takes electrical power from the line and uses that to regenerate a communications signal. The amount of power that can be supplied down the line limits the number of repeaters that can be supported on that line, thus limiting the actual deployment distance. Line powering regulations vary according to geography, but generally support between 2 and 4 repeaters from one power source depending on the local regulations. If power is sourced from both the central office and customer location, this dual-ended powering architecture effectively doubles the number of repeaters that can be supported on that line.

Since each repeater provides a brand new signal, the rate/reach of Mid-Band Ethernet from that regeneration point is the same as it is from a central office. However, as will be presented, actual repeater deployments generally operate at much less than full power for spectral compatibility reasons.

The Benefits of Regeneration…

The benefit of using repeaters to regenerate signals is reach. Using 2BASE-TL, 10 Mbps services can be delivered at distances beyond 11 Kft on 8-pairs, and to nearly 9 Kft on 4-pairs. A 5 Mbps Ethernet service can be delivered beyond 13 Kft on 8-pairs. To get service rates and reaches beyond these numbers, the signal-generation point must be moved closer to the customer.

The good news is that the vast majority of Mid-Band Ethernet applications do not require any repeaters. In North America, which has some of the longest loop lengths in the world, more than 75% of potential Mid-Band Ethernet customers would not require any repeaters because they are within the rate/reach of a non-repeatered 2BASE-TL solution. Another 15% of potential enterprise customers are at distances that may require at most a single repeater, depending on the required rate. In other countries, the situation is even better, with commonly 90-100% of customers having loop lengths that do not require repeater-supported deployments.

However, that final small percentage of customers that are at the longer distances are still a very important aspect of the customer base. In many cases, a carrier has to reach each and every customer site if they are to win that customer's business. Reaching 90% or even 95% of a given customer's locations may not be enough. In these select cases, repeatered deployments are a necessary and crucial part of winning the customer.

And The Difficulties

Repeatered deployments are rarely an attractive economic deployment strategy. The cost and complexity of deploying repeaters is high. All other strategies should be examined first, with repeater deployments being used as the final alternative.

The reason for the high cost and complexity is simple – it requires more stuff. As a simple example, assume there is a customer located 13 Kft from the CO that needs a 5 Mbps Ethernet service. This service could be delivered over 8-pairs without repeaters. Alternatively, it could be delivered over fewer pairs using repeaters. Using repeaters would require each line to be regenerated twice in that 13Kft (assuming a 6Kft deployment strategy typically used with traditional HDSL technologies). The repeatered solution requires multiple repeaters, line power equipment, and all of the difficulties that go along with remote line powered equipment. And of course this only applies to incumbent operators that own the outside plant – competitive carriers do not have access to the copper span and thus can never deploy repeaters. Using more pairs is always more economically attractive than using repeaters.

In some cases however, more pairs are not a viable alternative. A second alternative to be investigated is the use of a remote cabinet. Using a cabinet solution would allow multiple customers to be served, and not just a single customer as with a repeatered solution. With a remote cabinet, one piece of equipment provides service for many customers. With a repeatered solution, many pieces of equipment are needed to provide service for a single customer. If there are only one or two customers in that same vicinity, a repeatered approach may make more economic sense. However, if there are multiple customers that can be reached in that same vicinity, then a remotely-deployed cabinet will be more economically attractive than a repeatered solution.

So, repeatered deployments are in some sense a necessary evil. The next question to ask is how well repeaters can work in those few cases where they are absolutely necessary.

In an ideal world, each repeater would run at full power and full rate, thus regenerating a very high speed signal. Unfortunately, repeaters cannot realistically be deployed in that fashion due to spectral compatibility issues. Imagine having a conversation with someone, with someone else yelling at the top of their lungs inserted between the two of you. It would make your original conversation difficult, if not impossible. The same is true of putting a repeater running at full power in the midst of other electronic signals – the repeater's power would mask and corrupt other signals, thus rendering them incomprehensible.

Repeaters therefore, when used in the normal outside plant environment, generally operate at lower power and at lower rates. Studies have shown that a 2BASE-TL repeater can run at approximately 1.4 Mbps while maintaining spectral compatibility with other services. Thus Mid-Band Ethernet repeaters can run at rates higher than traditional T1/E1 repeaters (which run at 768 Kbps or 1024 Kbps), but not drastically higher. A repeatered application does not provide the full bandwidth benefits of a non-repeatered Mid-Band Ethernet deployment.

At the end of the day, carriers must make the most economic choice for service delivery. Non-repeatered alternatives need to be explored and exhausted before accepting the cost and complexity of a repeatered service, and will most often lead to more attractive options. But for that small percentage of customers where no other viable option exists, Mid-Band Ethernet services can be extended great distances by using mid-span repeaters.

2BASE-TL Versus Multiple T1/E1 Circuits

Carriers have multiple options for reaching the business customer with high-speed symmetric services. For example, business customers have traditionally been served with T1/E1 circuits. The natural integration of T1/E1 circuits into a SONET/SDH network, as well as the relative ubiquity of the service offering, has made T1/E1 circuits one of the most popular traditional choices for access technology. As the bandwidth needs of customers grow, it is natural to consider multiple T1/E1s as a higher speed access solution. This section looks at the differences between multiple T1/E1s and 2BASE-TL.

Access Architecture and Costs

Perhaps one of the most significant differences between a multi-T1/E1 solution and 2BASE-TL is on the access architecture and costs of the solution. 2BASE-TL requires plain copper access lines (also known as UNE-L in North America) for the customer service whereas multi-T1/E1 solutions require a T1/E1 interface (sometimes known as a UNE-P in North America).

Both solutions require the competitive carriers to lease something (UNE-L or UNE-P) from the incumbent carrier. The cost of these different access mechanisms is very dependent on geography. However, a UNE-L (dry copper pair) is generally much less expensive than a UNE-P. For example, in the USA, a UNE-L is available for $10-15/month and a UNE-P is available for $100-400/month. When using multiple lines for higher speed services, the cost difference is even more dramatic. To offer a 10 Mbps service using multi-T1/E1 lines would require 7 circuits and have a recurring cost to the provider of $700-$2800/month. To offer a 10 Mbps service using 2BASE-TL would require 2-8 pair (depending on the distance and noise environments) and have a recurring cost of $20-120/month. It's clear that the carrier offering a 10 Mbps service using 2BASE-TL could do so at significantly lower cost than the carrier offering that service using multiple T1/E1 lines.

Additionally, T1/E1s have in the past been universally available and regulated at reasonable prices. Recent rulings by the FCC in the USA, and by the regulatory control bodies in other countries, are changing that landscape. New rulings suggest that UNE-P based T1/E1s may not be available in the future, and if they are, they may come at unreasonably high prices. Many competitive carriers who have been serving the business market with T1/E1 access lines are being forced to find an alternative strategy to maintain a cost effective access solution going forward. 2BASE-TL offers a way around these regulatory changes while decreasing the carriers' costs.

More Bandwidth and New Services over the Same Copper Pairs

A multi-T1/E1 approach provides the carrier with a mechanism to deliver a higher bandwidth solution at a significantly higher cost and complexity than 2BASE-TL. The carrier can continue their traditional T1/E1 service offerings with simply a larger pipe to the customer. On the other hand, 2BASE-TL allows the definition of a new and differentiated service offering with significantly higher bandwidth and revenue potential. Instead of a 6 Mbps service over four T1 lines, 2BASE-TL can provide more than 22 Mbps over four access lines (and more than 45 Mbps over the 8 copper loops used to deliver those four T1s). The potential bandwidth differences allow a new and exciting range of services that would be difficult or impossible over legacy technologies such as T1/E1. In addition, 2BASE-TL equipment delivers plug-and-play provisioning, and is less costly in capital and operational costs than multi-link T1/E1 alternatives.

SONET/SDH or Packet

Additionally, 2BASE-TL and T1/E1 differ on the primary intended application. T1s and E1s were developed for voice transport and for integration into a SONET/SDH network. Although one can carry data as an overlay on a T1 or E1, just as one can carry voice as an overlay over a packet network, it is not the primary application for which the technology was designed. And as data continues to grow and completely dominate the network traffic, the network itself must focus on supporting that dominant application. Core and metro networks are already well underway in their migration to Ethernet, MPLS, and RPR. With 2BASE-TL, the evolution of the access network is now also underway.

As a pure Ethernet technology, 2BASE-TL provides a natural mechanism for next-generation services such as Ethernet VPNs, E-LINE, E-LAN, and VoIP. 2BASE-TL offers a copper access technology that can serve up the same services with the same, or better, guarantees as optical Gigabit Ethernet - the only difference is the available bandwidth. This compatibility with the optical packet core guarantees not only a consistent service offering to the end user, but also a simplified management strategy that reduces the overall cost of maintaining the network.

Dynamic Spectrum Management and Multiple Input Multiple Output

One of the key features of any loop based transmission system is the ability to efficiently utilize the available spectrum. Dynamic spectrum management (DSM) is a general term used to describe various techniques that permit the efficient use of spectrum. The techniques of DSM, though most often discussed in the context of VDSL2, are generally applicable to any copper loop technology.

DSM Background

DSM requires changes to standard xDSL in two ways. First, it requires more knowledge of the actual current status of the copper loops. Although xDSL technologies have always provided some information related to current status and performance of the lines, DSM requires much more detailed and granular information. Secondly, DSM requires additional parameters by which the xDSL technology can be finely tuned.

One of the most important examples of DSM is power back-off. Many systems today operate in a greedy mode of operation – they will transmit at full-power even if they can achieve their performance requirements using lower power transmission. Power back-off is basically the idea that the transmitter should not transmit excessive power – only take what you need.

Dynamic spectrum management requires a control system to probe the lines to retrieve the additional knowledge available about each loop, and then to go fine tune each loop using the knobs discussed earlier. The goal of all of this monitoring and fine tuning complexity is to increase the overall performance of the access network. In our power back-off example, DSM enables better power back-off by providing information to the control entity that there is excess power being used on a line, and allowing the control entity to adjust the transmit power down accordingly.

In certain cases, DSM has been shown to result in increased overall performance of the access network. In other cases, the performance benefit of DSM has been shown to be negligible. The inconsistent results have resulted in varying levels of support for DSM by both carriers and vendors. DSM, although it has great potential, does not always result in significant performance gains, and yet always carries significant complexity. To be truly beneficial, the benefits of DSM must be made more consistent, and the costs must continue to be reduced.

One aspect of DSM that is of questionable benefit is Multiple Input Multiple Output (MIMO). MIMO is a technique that attempts to eliminate cross-talk by coordinating transmission across multiple lines simultaneously, and is classified as one potential variant of DSM techniques. MIMO can be considered DSM to the extreme – it can have the most significant performance impact, but it also has that impact in only the most specific of circumstances.

Separating Reality From Hype

Although MIMO can improve performance, there are some incredibly important limitations and restrictions to recognize:

1) MIMO works best when all pairs are in the same binder with no alien disturbers. Unfortunately, carriers rarely know which pairs are in which binders, and it is impossible to guarantee that any 2 pair, even to the same destination, are in any way co-located in the same binder. And the problem gets exponentially more difficult if you try to coordinate across more than 2-pair. Additionally, MIMO doesn't perform nearly as well when there are alien (non-coordinated) disturbers. In the reality of the outside plant, where you can't predict which pairs are in the same binder or which pairs in the same binder exhibit cross-talk on one another to what degree, and where alien disturbers are certain, MIMO does not offer any guaranteed performance benefit.

2) MIMO doesn't behave well in dynamic environments. MIMO-based implementations do not handle changes well – they take down the entire connection to the customer. In bonded 2BASE-TL implementations, for example, a change may result in one pair going down and re-training, but the other pairs in the bonded group, as well as the bonded connection, always remains operational. With a MIMO-based technology, changes in the noise environment take down the entire bonded group. So deteriorating noise conditions (for example from the introduction of other xDSL loops into the binder, from wet/damaged pairs or connectors, etc.) or even just transient noise conditions (for example impulse noise, radio interference, etc.) can take down an entire bonded MIMO group. Likewise, changing the make-up of the bonded group, by adding or removing a pair, also results in the customer going DOWN. Multiple-pair operation must increase the resiliency of the

application to dynamic changes – MIMO has a hard time living up to that.

3) MIMO implementations are NOT VDSL2 and DO NOT have the same spectral properties. There is often confusion as to the distinction between MIMO and VDSL2. To be clear, VDSL2 has nothing to do with MIMO, and MIMO is completely different than VDSL2. The proprietary MIMO implementations available on the market today are not VDSL2 implementations. They are proprietary technologies that use a similar transmission mechanism to that of VDSL2 (and ADSL, and ADSL2, and ADSL2/plus, etc.). These proprietary MIMO technologies do not follow the same bandplans or profiles of VDSL2, will never interoperate with VDSL2, and do not have the same spectral characteristics as VDSL2. Proprietary MIMO implementations have not been ratified by standards bodies, and have not been certified for deployment in carrier networks. The use of non-standard, non-approved bandplans and technologies is a huge danger that must be avoided.

4) DMT-based MIMO doesn't perform well in business environments. As mentioned previously, the proprietary MIMO implementations on the market are based on DMT. Earlier in this paper it was shown how both ADSL2plus and VDSL2 have a difficult time operating in the business environment due to the very common presence of T1/E1 disturbers. When traditional T1/E1 technologies are present in the binder, DMT-based MIMO implementations suffer just as ADSL2/plus and VDSL2 suffered. These DMT-based systems, though excellent for residential applications where common business disturbers are not present, fail to deliver useful symmetric bandwidth for the typical business customer.

5) MIMO can't deliver universal reach. Unlike 2BASE-TL, MIMO cannot be deployed in a traditional repeater architecture. This is because MIMO requires centralized coordination and knowledge of the all of the loops and the cross-talk implications on all of the other loops in the same group. Traditional repeater architectures break this paradigm – they sit in the middle of the loops and have access to only one or two loops at most. This interruption in the line eliminates the possibility of knowing, let alone coordinating, cross-talk in a repeatered environment. As was discussed earlier, there is some percentage of business customers that sit on very long loops, at 20 Kft and beyond, and only

a repeatered technology, such as 2BASE_TL, can deliver 5 or 10 Mbps Ethernet services at those distances.

6) MIMO is just another modem function. MIMO has great potential to increase throughput in the right circumstances – it's a sexy technology. But at the end of the day, it's just another modem function. DSL chip vendors are already implementing cross-talk and MIMO functionality into the next generation xDSL silicon. As chip vendors integrate this function in cost and space effective architectures, some of the currently insurmountable obstacles in the way of useful MIMO deployments will be alleviated. And at the end of the day, MIMO - just like Reed-Solomon encoding, forward error correction, digital signal processing, and analog front ends - will be just another basic function available from any modem.

So MIMO is a promising technology, but it also has its problems – at least at present. Today it has a very high cost, and as we've discussed, it offers unpredictable performance results that depends on how the pairs are correlated, what alien disturbers are in the binder, and on whether T1/E1 disturbers are there or not. A technology with high cost and unpredictable results is almost impossible to accept for the enterprise environment, where carriers want to establish simple deployment guidelines that can quickly and reliably determine which customers can be reached with which services. With MIMO, you don't know what you get until you try.

MIMO implementations also currently lack the resiliency required for enterprise services, taking the entire subscriber connection down unnecessarily. Subscriber connectivity should never be interrupted, and in that regard MIMO implementations have significant problems. Add to that the high risk involved in such a proprietary, non-interoperable technology, and the fact that MIMO still has a more limited addressable market than 2BASE-TL, and the MIMO pill becomes increasingly difficult to swallow. But as international standards for MIMO emerge, and the silicon vendors integrate the technology into their modems, the interoperability and cost problems go away — finally resulting in a more practical application of science.

Key Requirements for Mid-Band Ethernet Solutions

Adhering to the IEEE 802.3ah 2BASE-TL standard is just the first requirement for providing a carrier-class Mid-Band Ethernet solution. Additional requirements must be met to reach the full potential of 2BASE-TL products. For example, full layer-2 switching with quality of service is necessary for carriers to truly differentiate their service offerings, and carrier-grade reliability, 5-9s, is expected in order to achieve the same reliability and dependability as today's TDM-based services. Furthermore, new deployments must achieve a fast return on investment (ROI), so a solution must have a low initial cost, but also be scalable.

Bandwidth in Copper Access Networks is Precious and Requires Intelligent Protection

The ability to deliver high-margin, differentiated services is fundamental to a carrier's ability to generate meaningful new revenue, to keep customers happy, and to compete for new business opportunities based on performance and quality rather than price. While some believe that an overabundance of bandwidth can serve as a means to achieve high quality of service, they are missing a key opportunity to leverage advances in Ethernet platforms capable of delivering carrier-grade Service Level Agreements (SLAs) based on a flexible service framework. And bandwidth, though sometimes abundant and inexpensive in the core, is still quite precious in the access network.

Quality of Service (QoS) technology provides a method for categorizing traffic and for ensuring that particular categories of traffic will always flow across the network at their specified bandwidth, latency, and jitter service levels, regardless of competing demands. For example, QoS can guarantee specific total network transit time (latency) and transit time variation (jitter) for particular traffic flows; the perceived quality of a VoIP service, for example, could be highly dependent on such guarantees. It can also provide guarantees of bandwidth, error rate, and many other characteristics.

For example, let's take a fully capable layer-2 Ethernet switch that supports two independent Quality of Service parameters for every packet that enters the system. The first parameter is the Class of Service (CoS), which is a priority level that controls the latency of the packet through the switch. The second parameter is the discard eligibility that controls the probability of the packet being discarded under congestion.

An ingress classification and policing process determines the priority and discard eligibility of each frame. Some Ethernet switches use the concept of a flow, or traffic stream, to control discard eligibility and determine how much traffic is allowed into and out of the network. In typical configurations, flows and traffic streams are associated with services or applications (e.g. Internet access, Transparent LAN Service, Voice over IP, etc.). A mapping of packet CoS markings (such as 802.1p-bits or IP DSCP bits) to CoS values can be configured for each port. Then, when provisioning services over that port, the CoS values and service identifiers are mapped into traffic flows (See Figure 8). These traffic flows are given bandwidth parameters, and the discard eligibility for each packet is based upon those parameters and the amount of traffic entering the system in that flow.

Also illustrated in Figure 9, the CoS Classification process gives the carrier the ability to normalize a subscriber's CoS markings into a space controlled by the service provider (the CoS identifier). Each customer can then be treated and served independently. The Flow Classification stage permits multiple traffic flows for each customer within each service. This allows, for example, different bandwidth controls for a VoIP application than for an email application from the same customer.

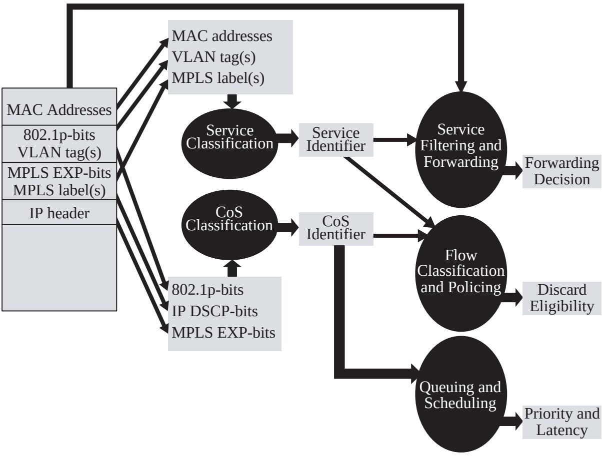

Figure 9. Determining Quality of Service

Packets are first classified into services to determine the forwarding rules for the packets, and then they are mapped into flows to control the priority and bandwidth of any data stream. Per flow discard eligibility parameters are used to treat traffic fairly during periods of congestion. The result is a solution that overcomes the limitations of ATM-based DSL access multiplexers (DSLAMs) and traditional Ethernet switches, with a QoS solution that integrates and interoperates for end-to-end service quality.

The Benefits of QoS in Mid-Band Ethernet Networks

SLAs, when provisioned on a per port and per service basis, can be used to tier services based on a flexible Quality of Service and Class of Service framework. All services can then be converged over a single physical connection with their own SLA.

Another advantage of using a QoS and CoS framework is the ability to help carriers deliver "stickier" services. Higher levels of customer satisfaction lead to "stickier" services, meaning business customers are less likely to seek alternative carriers for their data service requirements. Keeping customers happy, and loyal, is a key challenge for carriers today and in the future as they face renewed competitive threats from alternative carriers for higher speed data access, and from cable providers for lower speed access. Ethernet access can play a key role in maintaining customer loyalty by enabling carriers to deploy a wider range of services, delivering them more quickly and efficiently, and tailoring them to more closely match specific customer needs.

Today's Ethernet QoS and CoS capabilities offer substantial benefits to carriers and their customers alike. Improved application performance, along with support for new applications and services, such as point-to-point private line, point-to-multipoint, multipoint-to-multipoint, Transparent LAN, and Internet access, are all possible when QoS and CoS are deployed network-wide, beginning with the edge device.

Ensuring Complete Business-grade Resiliency

Business services require a level of resiliency beyond that of residential services. Resiliency, in order to be effective, must exist at every layer in the network hierarchy, as the service is only as resilient as its weakest link. Carriers today expect and demand new product offerings to provide carrier-grade availability of 5-9s. Today's IP/Ethernet products now carry time-sensitive traffic like voice and video, in addition to critical business applications. Achieving the same reliability and dependability as today's TDM-based services is a top requirement of carriers.

Mid-Band Ethernet products should be designed with true business-grade resiliency rather than a semi-redundant system. Resiliency is required across every aspect of the system and network. The following are necessary to provide the system and network redundancy that carriers need to satisfy the most demanding service up-time requirements: bonded copper ports to the subscriber; stackable nodes for system resiliency; spanning tree protocol and MPLS for network resiliency.

Key Features of a Business-grade Resilient Solution

- Full layer-2 Ethernet switch with spanning tree and rapid spanning tree protocols
- In multi-stack solutions:
  - A robust architecture preventing any single points of system failure using redundant interconnections with distributed management control agents
  - Distributed bonding of copper pairs across multiple units for ultimate business class resiliency, simplified pair management, and no stranded ports
- Complete carrier-grade OAM in the Ethernet network
- Controls to limit unicast and multicast bandwidth
- Controls per-subscriber MAC address resources

Economic Initial Deployment AND Scalable

With CAPEX constraints and intense competition among carriers, new product deployments must achieve a fast return on investment in order to pass the business case litmus test. Furthermore, the product must also be able to cost-effectively scale as “truck rolls” for upgrades is no longer a viable option.

In the past chassis-based products were known to be somewhat more scalable, yet more costly at first than stackable units. New advances in technologies and product architectures have made it possible to build a cost-effective and highly resilient stackable switch.

As new customers are added, additional switches can be inserted into the stack, to form a “virtual node,” which works and is managed as a single entity. A ring-based interconnection method links the individual units into a cohesive unit. An additional benefit of this architecture is that copper pairs can be distributed across several switches in the stack so that if one switch in the stack fails, the customer will lose some bandwidth, but not the entire service. Also, a carrier can provision two uplinks, using different switches in the rack, so that if one uplink fails, the traffic will continue uninterrupted as the other uplink/switch continues to function.

The initial upfront expenditure is just one part of the cost equation. Reducing ongoing operational costs will actually save a carrier more money in the long run. A fully-capable Ethernet switch provides plug-and-play turn up of services, reducing OPEX and time to revenue. In being a “true” Ethernet switch, interoperability with the installed base of Ethernet equipment is ensured, easing the process and cost of equipment and management integration.

Conclusions

To truly capitalize on the potential revenue and cost benefits of the next generation network, carriers must leverage a true packet infrastructure. For the foreseeable future, fiber has a severely limited footprint, which forces carriers to use the universal medium of copper as the predominant access method.

The Mid-Band Ethernet technologies of 2BASE-TL and 10PASS-TS were developed to allow carriers to utilize their existing copper infrastructure as high-speed, high-margin on-ramps to their metro and core packet networks. Native Ethernet in the access network lowers cost, simplifies deployments, and yields a more flexible network for new services such as VoIP. Like all Ethernet technologies, EFM has delivered an international interoperable standard with plug-and-play features as found with 10/100BASE-T. With higher-speed single pair performance, and with a new multi-pair aggregation paradigm that delivers speed and simplicity, Mid-Band Ethernet can revolutionize access architectures and services.

Adhering to the IEEE 802.3ah 2BASE-TL standard is just the first requirement for providing a carrier-class Mid-Band Ethernet solution. Additional requirements must be met to reach the full potential of 2BASE-TL solutions, including full layer-2 switching with quality of service, carrier grade reliability, a scalable architecture, and a low initial cost to achieve a fast return on investment (ROI).

With effective QoS and CoS capabilities, carriers can leverage Ethernet as a foundation for a new class of data services, layering multiple SLA-based services over each physical connection. These services, based on the performance and bandwidth requirements demanded by high-value business subscribers, will be fundamental to a carrier’s ability to cost-effectively compete for last-mile business customers on performance, not just price.

Hatteras Networks' Solution

Hatteras' Family of Metro Ethernet Service Edge™ Solutions

Figure 10. Family of Metro Ethernet Service Edge™ Solutions

The Metro Ethernet Service Edge™ product family enables service providers to deliver high-speed, symmetrical business-grade Ethernet services to businesses over their existing last-mile copper infrastructure. Hatteras Networks' Metro Ethernet Service Edge™ solutions allow service providers to fully leverage the benefits of offering data, voice over IP, and video services over Ethernet networks. With a low initial cost and plug-and-play installation, the product family easily scales from point-to-point applications to multipoint applications serving hundreds of customers from a single wiring center (Central Office, Remote Terminal, pole, or vault).

While other solutions offer transport and aggregation functions, the flexible architecture of the Hatteras Networks' solution provides the industry's first hardened, multi-service Ethernet system designed for the access network. Metro Ethernet Service Edge™ products allow multiple services per port using Hatteras Networks' patented VLAN product architecture to provide both E-Line and E-LAN services as defined by the Metro Ethernet Forum (MEF), and an innovative Hub and Spoke service. With this feature, Hatteras Networks' Mid-Band Ethernet solutions offer true Quality-of-Service, going beyond the simple Class-of-Service solutions of other vendors.

HN4000i Metro Ethernet Service Edge™ Switch

Hatteras Networks' HN4000*i* Metro Ethernet Service Edge™ switch is an innovative Ethernet in the First Mile (EFM) product that extends the reach of native Ethernet services to businesses that do not have access to fiber. The HN4000*i* delivers to each customer a 1-45 Mbps symmetric Ethernet service over 1 to 8 pairs of existing last-mile dry copper utilizing standards-based 2BASE-TL technology - effectively and efficiently bridging the existing T1/E1 – T3/STM-1 service gap and economic disparity. With the HN4000*i*, high-margin Ethernet services can now be delivered over voice-grade copper at full carrier serving area (CSA) distances and beyond.

Packaged in a stackable high-density 1RU "pizza box" the HN4000*i* is purpose-built for deployment throughout the access network – in a central office (CO), controlled environmental vault (CEV), or remote terminal (RT). The fully front-accessible platform is standards-based, temperature-hardened, highly resilient and scalable. The HN4000*i* supports an industry-leading density of 40 pairs of 2BASE-TL per rack unit.

The HN4000*i* switch is fully compatible with the Hatteras Networks' HN400-CP*i* series of cost-effective customer-premise demarcation devices. The connection between an HN4000*i* and an HN400-CP*i* device can consist of 1-to-8 copper pairs, bonded into a single 2BASE-TL broadband connection. IEEE 802.3ah 2BASE-TL copper pairs can be bonded from adjacent or non-adjacent binders, automatically bonding pairs into a logical connection when connected to an HN400-CP*i*.

The HN4000*i* implements the IEEE standard for Ethernet OAM with extensions for complete remote management and control to simplify deployment and management, while maintaining full interoperability with existing Ethernet switches, routers and Ethernet ADM interfaces.

In cases where maximum port density is required, Hatteras' patented Virtual Node (VN) technology enables service providers to deploy multiple load-sharing systems as a single managed entity. Unlike other vendor offerings, the Hatteras Networks' VN is managed as a single node, providing a fully-redundant architecture and industry leading port density.

HN4000 Virtual Node (VN) Metro Ethernet Service Edge™ Switch

Hatteras Networks' HN4000 Virtual Node consists of up to five HN4000*i*s (capable of supporting 200 2BASE-TL pairs), redundantly connected using a dedicated ring-based stacking interface to create a Virtual Node that enables carriers to incrementally add service capacity. HN4000*i* platforms can be hot-inserted into an operational stack, with no impact to existing services. The HN4000VN implements a distributed cross-box bonding mechanism that frees the carrier from the inventory and planning issues typically associated with multi-pair systems – any pair can be aggregated with any other pair, anywhere across multiple HN4000*i*s in the VN.

A fully-stacked HN4000VN supports all of the functionality of an HN4000*i* and is fully compatible with the Hatteras Networks' HN400-CP*i* series of cost-effective customer-demarcation devices.

HN4000i and HN4000VN Highlights

- Completely standards-based and spectrally compatible solutions
- Complete Quality-of-Service rather than simple class-of-service solutions
- Complete business-grade resiliency rather than a semi-redundant system
- Complete carrier-grade OAM on an Ethernet network - managed via industry-standard CLI, TL1, WebManager, EMS, SNMP, and IBM Tivoli Netcool®/OMNIbus integration
- Ultra-compact, front access products with industry leading port density per rack unit
- Full temperature hardening for versatile deployment
- NEBS3, ETSI and OSMINE certified
- Plug-and-play turn up of services reduces OPEX and time to revenue
- Patented VLAN-aware product architecture enabling multiple services per physical connection, with a smooth migration from tag stacking to MPLS
- Sharing of single or multiple network uplinks across hundreds of customers
- Distributed bonding of copper pairs across multiple units for ultimate business class resiliency, simplified pair management, and no stranded ports
- Robust architecture preventing any single points of system failure using redundant, Virtual Node interconnections with distributed management control agents

HN4000i and HN4000VN Applications

- High-speed symmetrical Ethernet transport/access
- Extension of existing fiber-based Ethernet services
  - Ethernet Private Line
  - Ethernet Virtual Private Line
  - Ethernet Private LAN
  - Ethernet Virtual Private LAN
  - Hub & Spoke

- Wireless tower and DSLAM backhaul
- MDU/MTU, campus networks, universities, government, etc.
- Converged access and transport where per-service SLAs are required for voice, video and data – all over a single physical connection

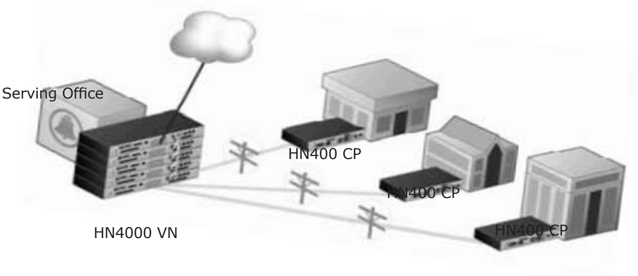

Figure 11. HN4000i Multipoint Application

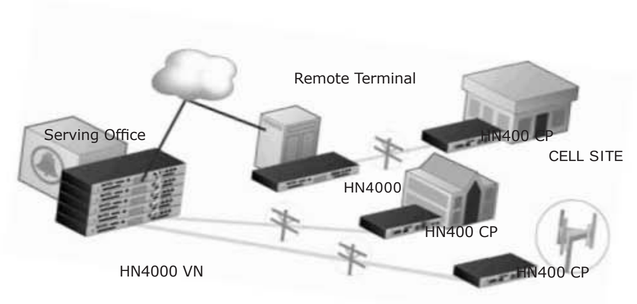

Figure 12. HN4000VN Multipoint Applications - CO and Remote Terminal

HN400-Ui Metro Ethernet Service Edge™ Switch

Hatteras Networks' HN400-U*i* (Universal for CO or CP deployment) is an innovative product that extends the reach of native Ethernet services to businesses without access to fiber. The HN400-U*i* delivers point-to-point symmetrical Ethernet bandwidth and services at rates up to 45 Mbps over existing copper pairs utilizing standards-based 2BASE-TL technology. The HN400-U*i* provides a low-cost market-entry solution for extending Ethernet to the customer premise.

The HN400-U*i* is a carrier-class, temperature-hardened platform that enables carriers and service providers of all sizes to deliver service-rich business broadband services on basic last-mile copper pairs. By extending native Ethernet to business customers without fiber access, Hatteras Networks' HN400-U*i* offers greater service flexibility and efficiency than existing private-line or asymmetrical services do today. Using the HN400-U*i*, carriers are able to leverage existing copper assets to deliver new services with the quality of traditional private-line services.

The HN400-U*i* platform is fully compatible with the Hatteras Networks' HN400-CP*i* series of cost-effective customer-demarcation devices. The connection between an HN400-U*i* and an HN400-CP*i* device can consist of 1-to-8 copper pairs, bonded into a single 2BASE-TL broadband connection.

The HN400-U*i* implements the IEEE 802.3ah standard for Ethernet OAM with extensions for complete remote management and control to simplify deployment and management. The HN400-U*i* uses essentially the same CLI as the HN4000*i*.

Hatteras Networks' HN400-U*i* reduces the risk and enhances the reward associated with offering new services by leveraging the existing copper infrastructure and seamlessly integrating into the carrier's operational processes.

HN400-Ui Highlights

- High-speed symmetrical bandwidth and Ethernet services over copper
- Cost-effective point-to-point solution - profitable with a single customer
- Standards-based utilizing 802.3ah - 2BASE-TL
- Full temperature hardening for versatile deployment
- NEBS3, ETSI and OSMINE certified
- Complete carrier-grade OAM on an Ethernet network - managed via industry-standard CLI, TL1, WebManager, EMS, SNMP, and IBM Tivoli Netcool® /OMNIbus integration
- Plug-and-play turn up of services reduces OPEX and time to revenue
- Easy migration to HN4000*i*s as more customers are added to the network

HN400-Ui Applications

- High-speed symmetrical Ethernet transport/access
- Ethernet Private Line Service
- Ethernet Virtual Private Line Service
- Wireless tower and DSLAM backhaul
- MDU/MTU, campus networks, universities, government, etc.

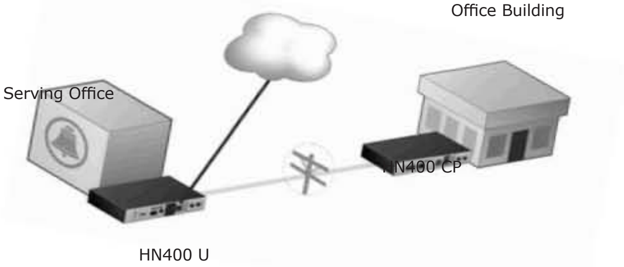

Figure 13. HN400i Point-to-Point Application

About Hatteras Networks

Hatteras Networks is headquartered in Research Triangle Park, North Carolina and was started several years ago to develop high-performance telecommunications access platforms. Hatteras was formed to enable telecommunications service providers to deliver cost-effective, service-rich, high-speed access in the last mile to business customers - one of the most difficult challenges carriers face today. Removing this critical bottleneck in the access network is necessary to fuel future growth in broadband, multimedia and Internet services. To date, capital-intensive fiber-based solutions have been the most widely deployed means for meeting this challenge.

Hatteras Networks provides standards-based Ethernet access solutions, which leverage the fully ratified Ethernet in the First Mile (EFM) standards from the IEEE and ITU. With Hatteras Networks' solutions, carriers can migrate from the complexity and expense of TDM-based T1/E1 circuits, to the simplicity and availability of a pure Ethernet access platform, all while increasing the bandwidth to, and revenue potential of each customer.

Hatteras is a founding member of the Ethernet in the First Mile Alliance, and was a leader in the development of the IEEE 802.3ah standards for delivering symmetrical Ethernet services natively over copper access loops. Hatteras' corporate commitment to standards and interoperability provides our customers with the certainty that their capital investments are protected.

Carriers from around the world have deployed Hatteras' Mid-Band Ethernet solutions, enabling them to drive down access costs by eliminating the cost and complexity of ATM and T1/E1 solutions, and increase revenue with higher bandwidth, value-added services.

Visit www.hatterasnetworks.com for additional information.

Extending the Ethernet Service Edge

Hatteras Networks enables an emerging market segment referred to as Mid-Band Ethernet. The Mid-Band Ethernet service enables Carriers to deliver the Metro Ethernet services over the existing copper infrastructure to businesses whose application requirements fall within the bandwidth gap between T1/E1 and T3/STM-1. While the Hatteras solutions deliver up to 45 Mbps over 8 copper pairs, the Mid-Band Ethernet service sweet spot is 2-45 Mbps (the gap between T1/E1 and where fiber deployment becomes economically viable).

Mid-Band Ethernet services are exactly the same as those enabled by Metro Ethernet:

- Transparent LAN Services (TLS)
- Direct Internet Access (DIA)
- Voice over IP (VoIP)
- Ethernet Private Line
- Storage Area Networks (SANs)
- etc…

Therefore, an essential requirement of Mid-Band Ethernet equipment is the ability to transparently extend existing Metro Ethernet services beyond fiber.

Vertical System Group estimates that in the U.S. and Europe over 2.2 million T1/E1, Frame Relay and T3/STM-1 connections will migrate to Mid-Band Ethernet links over the next 5 years. Hatteras Ethernet Service Edge solutions enable this new market segment that is expected to generate over $15B per year in service revenue for U.S. and European Carriers. The opportunity in other international markets is likewise compelling.

HATTERAS™
NETWORKS

P.O. Box 110025

Research Triangle Park, NC 27709-0025

TEL: 1-919-991-5440

FAX: 1-919-991-0731

sales@hatterasnetworks.com

www.hatterasnetworks.com